

APPLICATION
FOR
UNITED STATES LETTERS PATENT

TITLE: RADIO NETWORK CONTROL

APPLICANT: M. VEDAT EYUBOGLU, ARTHUR J. BARABELL AND
SANJAY CHERIAN

CERTIFICATE OF MAILING BY EXPRESS MAIL

Express Mail Label No. EL298426440US

I hereby certify under 37 CFR §1.10 that this correspondence is being deposited with the United States Postal Service as Express Mail Post Office to Addressee with sufficient postage on the date indicated below and is addressed to the Commissioner for Patents, Washington, D.C. 20231.

June 25, 2001
Date of Deposit

Joshua Cronin
Signature

Joshua Cronin
Typed or Printed Name of Person Signing Certificate

RADIO NETWORK CONTROL

BACKGROUND

This invention relates to radio network control.

High Data Rate (HDR) is an emerging mobile wireless access
5 technology that enables personal broadband Internet services to be
accessed anywhere, anytime (see P. Bender, et al., "CDMA/HDR:
A Bandwidth-Efficient High-Speed Wireless Data Service for
Nomadic Users", IEEE Communications Magazine, July 2000, and
3GPP2, "Draft Baseline Text for 1xEV-DO," August 21, 2000).
10 Developed by Qualcomm, HDR is an air interface optimized for IP
packet data services that can deliver a shared forward link
transmission rate of up to 2.46 Mbit/s per sector using only (1X)
1.25 MHz of spectrum. Compatible with CDMA2000 radio access
(TIA/EIA/IS-2001, "Interoperability Specification (IOS) for
15 CDMA2000 Network Access Interfaces," May 2000) and wireless
IP network interfaces (TIA/EIA/TSB-115, "Wireless IP
Architecture Based on IETF Protocols," June 6, 2000, and
TIA/EIA/IS-835, "Wireless IP Network Standard," *3rd Generation
Partnership Project 2 (3GPP2)*, Version 1.0, July 14, 2000), HDR
20 networks can be built entirely on IP technologies, all the way from
the mobile Access Terminal (AT) to the global Internet, thus
taking full advantage of the scalability, redundancy and low-cost of
IP networks.

HDR has been adopted by TIA (Telecommunications Industry
25 Association) as a new standard in the CDMA2000 family, an
EVolution of the current 1xRTT standard for high-speed data-only
(DO) services, formally referred to as 1xEV-DO or IS-856.

IS-856 systems are typically implemented using the radio access network architecture shown in Figure 1. Here the Access Terminal (AT) 10 may be a laptop computer, a Personal Digital Assistant (PDA), a dual-mode voice/data handset, or another device, with
5 built-in IS-856 support.

The entire administrative service area of a wireless access provider may be divided into one or more subnetworks (or subnets) 12, 14. Each subnet 12 includes a set of Radio Nodes (RN's) 16, 18 and one or more Radio Network Controllers (RNC) 20, 22. The RN's
10 are connected to RNC's over a backhaul network 24. In existing 2G and 3G wireless networks, each RN is connected to only 1 RNC using dedicated leased lines or ATM permanent virtual circuits (PVC's). Further, RNC's are connected to each other using dedicated leased lines or ATM PVC's. In a new generation of IP-
15 based radio access networks, the backhaul can be implemented using a shared IP or metropolitan Ethernet network which supports many-to-many connectivity between RN's and RNC's.

Each RNC typically controls 25-100 RN's. Each RN typically supports 1-4 carriers each of 1.25 MHz of bandwidth. Further,
20 each cell area (not shown) is typically divided into multiple sectors (typically 3 or 6) and the RN has one radio transceiver 27 for each sector.

Each RNC is connected over an IP network 26 to one or more Packet Data Serving Node's (PDSN's) 28 (see TIA references
25 cited earlier). The RNC communicates with the PDSN over a standard interface termed the R-P (Radio-Packet) interface 30.

Attorney Docket 12144-007001

The R-P interface is further broken into 2 interfaces: the A10 interface used to carry data and the A11 interface used to carry signaling. A PDSN can be viewed as an edge router that supports mobility; it maintains link layer connectivity to AT's through the
5 Access Network. The PDSN also interfaces to AAA servers 32 for Authentication, Authorization and Accounting (AAA).

In IS-856 radio access networks as currently defined by 3GPP2 in 1xEV-DO IOS Phase 1 (IS-878), each RN is uniquely associated with an RNC and each subnet contains only one RNC. As a result
10 when an AT moves from the coverage area of one RNC to the coverage area of another, it performs a handoff.

Every time a dormant AT crosses a subnet boundary, it initiates a dormant handoff by sending a UATI_Request. The AT recognizes the need for a dormant handoff by monitoring the 128-bit SectorID
15 being broadcast by the sectors. All sectors that belong to the same subnet have SectorID's that fall within a certain range. The 128-bit Universal Access Terminal Identifier (UATI) assigned to an AT in a given subnet falls within the same range. When the AT moves into the coverage area of another subnet, it compares its UATI with
20 the SectorID being broadcast by its serving sector. When these do not belong to the same range, the AT knows that it has crossed a subnet boundary and initiates the dormant handoff by sending a UATI_Request.

A first purpose of a dormant handoff is to inform the PDSN to
25 send packets arriving for that AT to the new serving RNC. Dormant handoffs involve a relocation of the R-P (A10) session

from the old serving RNC to the new serving RNC. Without such handoffs, the PDSN would send packets to an old serving RNC. Since the old serving RNC does not know the location of the AT outside its subnet, AT's packets may be lost.

- 5 A second purpose of a dormant handoff is to transfer session information between RNC's. In IS-856, each RNC maintains certain session information about the AT. Such session information is needed for communication over the air interface. Session information includes the Universal Access Terminal
- 10 Identifier (UATI), security keys for access channel authentication and encryption, and other protocol constants. Everytime the AT crosses an RNC boundary (in this case a subnet), a new UATI needs to be assigned to the AT and the remaining session information needs to be transferred from the old serving RNC to
- 15 the new serving RNC. Such a transfer requires a network link between the RNC's. Without such session transfer, every handoff between RNC's would result in a new and lengthy session establishment, taking up precious air resources and causing delays. When the footprint of an RNC is small, dormant handoffs occur
- 20 frequently, resulting in excessive use of airlink resources (for the new UATI assignment), extra processing for the RNC's to implement the session transfer, and extra processing for the RNC and PDSN to relocate the A10 connection.

- To address a similar situation in CDMA2000 networks, a new
- 25 logical network element often referred to as a Packet Control Function (PCF) has been added. The PCF provides a single-interface to a PDSN and serves multiple RNC's. Therefore, when

Attorney Docket 12144-007001

the AT crosses an RNC boundary, an A10 relocation is not required as explained below.

In CDMA2000 systems, mobility management is provided by the Mobile Switching Center (MSC) / Visitor Location Register (VLR). When an AT provides a location update to inform the network of its new location, this information is immediately forwarded to the serving MSC. Such location updates are provided by the AT when it moves sufficiently away from the sector where it last provided a location update. When the PDSN receives a packet data for an AT, it sends the packet to the PCF. The PCF, recognizing that no traffic channel exists for this AT, in turn informs the RNC that last served this AT. That RNC then informs the MSC. The MSC, equipped with location information stored in the VLR, requests one or more RNC's to page the AT from a certain set of sectors. Once the AT responds with an Origination message to its serving RNC, the serving RNC sets up a so-called A8/A9 connection to the PCF. Soon after, the PCF starts forwarding the received packets to the serving RNC.

A similar procedure can be used in IS-856, by adding a similar PCF entity 50, as shown in Figure 2. (A system like the one described here has been proposed in 3GPP2 for 1xEV-DO IOS Phase 2.) However, since IS-856 networks do not have an MSC, it is necessary to define a separate mobility management function (Mobility Manager 52) responsible for maintaining the location information for every AT that is being served by the RNC's under its control. Such a Mobility Manager can be integrated into the PCF, or can be a separate network entity. Any time the AT

Attorney Docket 12144-007001

provides a location update by sending an IS-856 RouteUpdate message, its location information is immediately forwarded by the serving RNC to the Mobility Manager. When the PDSN receives a packet data for an AT, it sends the packet to the PCF, which in turn
5 informs the Mobility Manager. The Mobility Manager, equipped with the location information requests, either directly or via the RNC that last received a location update from the AT, requests one or more RNC's to page the AT from a select set of sectors. Once the AT responds with a ConnectionRequest message, the serving
10 RNC sets up a so-called A8/A9 connection to the PCF. Soon after, the PCF starts forwarding received packets to the serving RNC. This approach eliminates the need for relocating the A10 (R-P) session to the PDSN every time the AT crosses an RNC boundary, effectively expanding the size of the subnet to cover multiple
15 RNC's.

The Mobility Manager function does not address the session transfer issue described earlier. For this purpose, another logical network element, a Session Manager 53 is introduced. Like the Mobility Manager, a Session Manager controls multiple RNC's,
20 and maintains session information for all sessions handled by the RNC's that it controls. Like the Mobility Manager, the Session Manager may be a separate network element, may be combined with the Mobility Manager or may be integrated with the PCF.

When a new session is to be established, the serving RNC interacts
25 with the Session Manager. The Session Manager provides the UATI to be assigned to the AT and stores the session parameters that the serving RNC has determined during the key exchange and

Attorney Docket 12144-007001

configuration phases of the session set-up. Whenever the AT establishes a new connection with its serving RNC, the RNC retrieves the session information from the Session Manager. In the case where the Session Manager is integrated with the PCF, this can be accomplished during the A8/A9 connection set-up procedures. The RNC provides the latest session information back to the Session Manager when a connection is closed. Again, in the case where the Session Manager is integrated with the PCF, this can be accomplished during the A8/A9 connection tear-down procedures. Additional delays are caused by passing of session information back and forth between the Session Manager and the RNC during every connection set-up.

The proposed 1xEV-DO Phase 2 architecture also produces frequent inter-RNC soft handoffs. This usually requires complex procedures that involve the old and new serving RNC's.

SUMMARY

In general, in one aspect, the invention features (a) in connection with a mobile wireless subnetwork including multiple radio network controllers and multiple radio nodes, associating a session established for an access terminal with a serving radio network controller, (b) maintaining the association as the access terminal moves from the coverage area of one radio node to the coverage area of another radio node within the same subnetwork, and (c) routing access channel packets from an access terminal having an existing session to the serving radio network controller by

determining the IP address of the serving radio network controller using a session identifier.

Implementations of the invention may include the following features. The routing is performed by an RN or by a broker radio network controller in the subnetwork. An RN forwards a received access channel packet to the broker radio network controller. The serving radio network controller and the broker radio network controller are connected by a high-speed LAN. The serving radio network controller and the broker radio network controller are connected by a high-speed LAN. The session identifier includes the Universal Access Terminal Identifier (UATI) of the IS-standard. The radio node routes packets received from an access terminal without an existing session to a default RNC with whom the radio node is associated. A radio node receives paging requests from more than one radio network controller and forward link traffic channel packets from more than one radio network controller. A radio node sends reverse link traffic channel packets to more than one radio network controller. Traffic channel radio resources are managed in the radio nodes and a radio network controller requests radio resources from a radio node before adding any of its sectors to a traffic channel. The radio network controllers reside in different locations and are connected via a metropolitan-area network. The session association is transferred from one radio network controller in one subnetwork to another radio network controller in another subnetwork based upon a predetermined criterion. The session transfer is triggered by the access terminal upon detection of a subnet change or by the

network. At the serving radio network controller, a packet data serving node is selected to serve the access terminal. At the serving radio network controller, a mobility manager maintains a current position of the access terminal.

- 5 An RNC Resource Control Agent assigns sessions to radio network controllers. The RNC Resource Control Agent resides on a separate server. An RNC Resource Control Agent also determines the association between the RN's and their default RNC's. The RNC Resource Control Agent also performs load
- 10 balancing in assigning sessions to radio network controllers. The RNC Resource Control Agent also selects a new RNC in network-initiated dormant handoffs. The Radio Resource Control Agent function is distributed among the radio network controllers and radio nodes, and the radio network controllers and the radio nodes
- 15 continuously communicate resource information to each other to enable individual network nodes to make session assignment decisions on their own. The Radio Resource Control Agent also maintains session information for all sessions under its control.

The radio network controllers also include a PDSN function that

20 includes the Simple IP, Mobile IP and AAA client functions.

In general, in another aspect, the invention features the radio node configured to route access channel packets from an access terminal having an existing session to a serving radio network controller by determining the IP address of the serving radio network controller

25 using a session identifier. Other advantages and features will

become apparent from the following description, and from the claims.

DESCRIPTION

Figures 1 through 4 show networks.

- 5 Existing 3G wireless network architectures, including the ones discussed above for IS-856, assume a fixed association between RN's and RNC's. In other words, all traffic flowing from or to an RN always goes through the same RNC. This requires complex hierarchical structures to handle dormant handoffs between RNC's and requires frequent and delay-prone inter-RNC (soft) handoffs. Fixed associations between RN's and RNC's are needed in circuit-switched voice applications when point-to-point dedicated leased lines are used for backhaul connectivity between RN's and RNC's as illustrated in Figures 1 and 2.

15 CLIENT/SERVER ARCHITECTURE FOR RNC CLUSTERS

- First, as shown in Figure 3, consider the case where a set of RNC's are co-located in a data center and are connected together via a high-speed Local Area Network (LAN) 62 such as a Gigabit Ethernet LAN. In this case, RNC's connect to the network via LAN interfaces and a router 64 provides the connectivity to the external network. We refer to such a configuration as an RNC cluster (or pool). (Later, we will describe how the same concept can be extended to RNC's connected over a metropolitan-area network.) Such clustering is possible for packet data and Voice over IP (VoIP) applications. In the past, when the main traffic type

Attorney Docket 12144-007001

carried through a wireless network was circuit-switched voice, such clustering using an Ethernet LAN was not feasible. RN's may connect to the router in the data center using dedicated leased lines 66. We assume that RN's and RNC's are all IP addressable.

- 5 In other words, any RN served by the cluster can communicate directly at the IP level with any of the other RNC's in the cluster. But in this section we assume that no direct communication link exists between RN's that are served by a cluster and other RNC's outside the cluster.
- 10 When constructing an RNC cluster such as the one described above, it is important to avoid any handoff boundaries between individual RNC's so that the entire cluster can behave as if it were one big RNC. This would eliminate unnecessary handoffs due to mobility, thereby greatly improving scalability and reliability.
- 15 To accomplish this, suppose we define an IS-856 subnet 70 to be the entire footprint of the RNC cluster, not the footprint of just one RNC. In other words, all the RN's served by the cluster now belong to the same subnet. To simplify the system operation, we continue to associate each RN in the subnet with only one RNC in
- 20 the cluster. This association is established when an RN is first powered. The detailed meaning of this association will be explained later.

Access Channel Packet Routing

- Each sector in an RN can transmit to an AT over the forward
- 25 traffic or control channels 72. Similarly, each sector in an RN can receive from an AT over the reverse traffic or access channels 74.

Attorney Docket 12144-007001

The access channel and the reverse traffic channels are separated by code division multiplexing using a Long Code Mask, whereas the control channel and the forward traffic channels are separated by time-division multiplexing using a preamble. The preamble
5 identifies a forward link physical layer packet as a control channel packet or as a traffic channel packet associated with a certain MAC Index. A MAC Index, an integer between 0 and 63, is unique within a sector and is assigned by the RN and RNC at the time of connection establishment. Similarly, the Long Code Mask
10 identifies a reverse link physical layer packet as an access channel packet or a specific traffic channel packet. The Long Code Mask is based on the AT's UATI for the traffic channel, and is based on the SectorID of the serving sector for the access channel. The sending AT of an access channel packet and the recipient AT of a
15 control channel packet are indicated in the ATI field of a MAC Layer header.

Whenever an RN receives a MAC Layer packet on one of its access channels, it always forwards the packet, without even looking at its content, to its Default RNC in the cluster with whom
20 it is associated. As such, when a packet carrying a UATI_Request message is received from an AT, it is forwarded by the receiving RN to the Default RNC. The RN encapsulates the MAC Layer packet in an IP packet (possibly multiplexed with MAC Layer packets of other AT's) with a destination IP address equal to an IP
25 address of the serving RNC. The IP packet is carried over the backhaul network to an aggregation router at the data center and the router forwards it to the serving RNC over the Ethernet LAN.

Attorney Docket 12144-007001

All access channel packets include an address field that identifies the sending AT. When the sending AT has already been assigned a UATI by the subnet, the address field contains that UATI. When the sending AT does not yet have a UATI, the address field

5 contains a Random Access Terminal Identifier (RATI), which is randomly selected by the AT. The first two bits of the address field indicate whether the address is a UATI or a RATI.

When the (Ethernet) I/O subsystem of an RNC receives a UATI_Request message from an AT with an address field that

10 contains a RATI or an unrecognized UATI, the RNC assumes the role of the serving RNC to handle the session and assigns the session to one of its server cards. The AT is then assigned a UATI within some predetermined range. This range, which identifies the serving RNC to all other RNC's in the cluster, is known by all the

15 RNC's in the cluster, but is not known by the AT. The range of the UATI's that belong to a certain RNC may further be subdivided to identify the server module within the serving RNC that is handling the session. The serving RNC also establishes an A10 connection with the PDSN in order to facilitate the data transfer between the

20 AT and the PDSN. The A10 connection terminates on the server module handling the session.

Page Routing

While dormant, the AT sends RouteUpdate messages, as needed, to provide information about its current location. This mobility

25 information is maintained at a Mobility Manager in the serving RNC. Since a subnet covers the entire footprint of the RNC cluster,

Attorney Docket 12144-007001

when the AT crosses the boundary between two RNC's in the same cluster, it does not detect a subnet change and therefore does not initiate a dormant handoff. But when the AT sends an access channel message to an RN that is associated with a different RNC

5 (broker RNC) in the cluster, the packet(s) carrying that message are sent by the RN to the broker RNC. The I/O subsystem in the broker RNC examines the address field of all arriving access channel packets and reads the UATI. From the UATI, the I/O subsystem determines by table look-up the identity of the serving

10 RNC and forwards the access channel packet to that RNC over the high-speed LAN. When a UATI is served locally, the I/O subsystem first determines the server module that is handling the session and forwards the packet to that module using an internal bus of the serving RNC.

15 If packet data is received from the PDSN for a dormant AT, the packets are always forwarded over the A10 interface to a specific server module on the serving RNC. That server module then obtains the location information for that AT from the Mobility Manager in the serving RNC. The serving RNC then sends a

20 paging message via a set of RN's that are determined based on the last Route Update message received from the AT. The paging message is sent via the control channel of one or more sectors that belong to the RNC cluster. The RN's transmitting the paging message may not be associated with the serving RNC (i.e., they

25 may have a different Default RNC), but they need to be associated with one of the RNC's in the cluster.

Connection Establishment

Attorney Docket 12144-007001

When it receives a ConnectionRequest message from the AT,
either directly or via a broker RNC, the server module in the
serving RNC examines the pilot strengths reported by the AT in
the RouteUpdate message accompanying the ConnectionRequest
5 message. To simplify system operation, we assume that each RN's
radio resources are managed by a Radio Resource Control function
in the RNC with whom the RN is associated. Therefore, when the
serving RNC wants to establish a connection that involves RN's
that are associated with other RNC's, it first communicates directly
10 with the Radio Resource Control function on those RNC's to check
for resource availability. Such communication occurs over the
high-speed LAN. When sufficient radio resources are available,
the serving RNC establishes the necessary traffic channel
communication links with the RN's and sends a
15 TrafficChannelAssignment message to the AT to initiate the
connection set up. Once a traffic channel has been established
packets flow directly between the RN's and the serving RNC
without any involvement of any broker RNC. Such direct routing
eliminates the delays typically found in soft handoff procedures
20 that involve triangular routing through another RNC.

When a new connection involves an RN that is outside the
footprint of the RNC cluster (different subnet), a soft handoff
procedure is implemented. In this case, the serving RNC
communicates with the RNC's outside the cluster over a
25 metropolitan-area network to obtain radio resources. If the radio
resources are available, the serving RNC establishes a
communication link with that RN, but this time through the RNC

outside the cluster. Such triangular routing is needed because here we assumed that there is no direct physical link between the serving RNC and the RN outside the subnet.

Dormant Handoffs/Subnet Change

- 5 When the AT crosses the boundary of an RNC cluster, it will detect a subnet change and initiate a dormant handoff between its serving RNC in the cluster and a new RNC outside the cluster. This handoff involves the assignment of a new UATI by the new RNC, the transfer of the IS-856 session from the old RNC to the
10 new RNC and the relocation of the A10 interface from the old RNC to the new RNC.

IMPROVED CLIENT/SERVER ARCHITECTURE

- The scheme described so far can be improved in couple of areas. First, we can eliminate the triangular routing of access channel
15 packets via a broker RNC by moving that routing function to the RN's. This will reduce delays in handling access channel packets, for example during connection set-up, at the expense of some increase in processing power at the RN. Eliminating triangular routing will also allow us to extend some of the benefits of RNC
20 clusters to RNC's that are connected across a metropolitan-area network.

- Second, as shown in figure 4, for backhaul networks 80 that support many-to-many connectivity between RN's and RNC's, two additional improvements are possible: A) We can extend a subnet
25 beyond the boundaries of a single RNC cluster, by allowing the

Attorney Docket 12144-007001

serving RNC to send Page Requests to RN's who are associated with an RNC that is not in the same cluster; B) We can move the Radio Resource Control function from the RNC's to the RN's, thereby further reducing delays in connection set-up procedures.

- 5 These improvements allow us better to exploit the flexibilities of IP and metropolitan Ethernet networks and result in a more distributed system where an AT may remain attached to its serving RNC regardless of its position, except when the distance between the AT and the serving RNC becomes excessive.

10 **Avoiding Triangular Routing of Access Channel Packets**

- When powered on for the first time, an AT registers with the IS-856 network as follows: It acquires an IS-856 pilot being broadcast by one of the nearby sectors and synchronizes with the system. To initiate the session establishment, the AT sends a
- 15 UATI_Request. As before, the AT uses a Random ATI (RATI) in the MAC Layer header to send this request.

- The RN examines the address field of the access channel packet and recognizes that the originator of the message does not have an assigned UATI and forwards the packet to its Default RNC with
- 20 whom it is associated, possibly when the RN is first installed. To examine the address field, the RN first extracts the MAC Layer capsule fragments from the received MAC Layer packets, and forms the MAC Layer capsule. It then reads the address field in the MAC Layer header.

After receiving the UATI_Request, the Default RNC assumes the role of the serving RNC and assigns a UATI to the AT. It then proceeds with the rest of session establishment, in particular the security key exchange and the protocol configurations. (Later, we
5 will describe an improved version of this procedure to increase availability and to provide better load balancing.) The RNC also implements the PPP/CHAP procedure to authenticate the AT based on its Network Access Identifier (NAI). There is a one-to-one mapping between the NAI and the terminal's actual IMSI
10 (International Mobile Subscriber Identity). This mapping is maintained in a AAA (Radius) server (not shown). The AAA server passes the AT's IMSI value to the serving RNC.

The PCF function in the serving RNC uses this IMSI value to select a PDSN as described in the IS-2001 standard and establishes
15 an A10 connection to that PDSN. In the A11 Registration message, the PCF function provides the IMSI value of the AT along with its own SID/NID/PZID identifier to the PDSN. The AT and the PDSN then set up a PPP link, perform Simple IP or Mobile IP set-up and execute user-level authentication.

20 Each RN keeps a routing table for the mapping between the UATI and the serving RNC. This routing table may be provided to the RN by a network management system. As in the previous system, each RNC owns the UATI values that fall within a certain range. Whenever the RN receives an Access Channel packet, it
25 determines from the UATI value in the MAC Layer Header the identity of the serving RNC, and routes the packet to that RNC by placing an IP address of the serving RNC in the destination address

Attorney Docket 12144-007001

field of the IP header. This approach allows access channel packets to be delivered from any RN directly to the serving RNC in the cluster.

Avoiding Handoffs Between RNC's Which Are Not Co-

5 Located

Now suppose we replace the point-to-point lines between RN's and RNC clusters, by a many-to-many backhaul network that allows any RN to communicate directly with any RNC, regardless of the location of the RNC's. In such networks we will of course benefit
10 from the direct routing of access channel packets from the RN to the serving RNC, by avoiding the triangular routing across the metropolitan-area network.

As before, mobility management for a given AT is handled entirely by the serving RNC. The AT is configured to provide distance-
15 based location update in dormant mode. In other words, whenever the serving sector is more than a certain distance away from the sector where it last sent a RouteUpdate message, the AT sends a new RouteUpdate message to the serving sector over the Access Channel. The RouteUpdate message is forwarded by the RN to the
20 serving RNC which then keeps track of the location of the AT.

When the serving RNC wants to page an AT, it first determines the RN or RN's from which it wants to send the page, depending on the time and position indicated in the most recent RouteUpdate message received from the AT. It is assumed here that the serving
25 RNC knows the IP addresses of all the RN's in the radio access network. The serving RNC sends the paging message to the

Attorney Docket 12144-007001

appropriate set of RN's directly. These RN's then page the AT over their respective control channels.

All sectors in an IS-856 network broadcast in their overhead channel their SectorID and Subnet Mask. For a relatively small
5 network, one can set the Subnet Mask to 0, thereby implying that the entire network is one big subnet. In this scenario, the AT never detects a subnet change. Therefore, the AT remains attached to the original serving RNC and never triggers a dormant inter-RNC handoff. The A10 connection to the PDSN also remains fixed
10 regardless of the position of the AT.

If the radio access network covers a geographically large area, it may be prudent to force a dormant inter-RNC handoff, when the AT moves too far away from the serving RNC. In this case, the Subnet Mask should be chosen greater than 0. Then, when the AT
15 crosses the subnet boundary, a dormant handoff occurs and the A10 connection is relocated. Further, the AT is assigned a new UATI and session parameters are transferred from the old serving RNC to the new serving RNC.

Faster Connection Using Distributed Radio Resource Control

20 Now we describe how moving the Radio Resource Control from the RNC's to the RN's reduces the set-up time for connections that involve multiple RNC's. Whenever the AT sends a ConnectionRequest message over the access channel along with a RouteUpdate message to initiate a new connection, the message is
25 immediately forwarded from the receiving RN to the serving RNC. The serving RNC examines the RouteUpdate message to determine

Attorney Docket 12144-007001

a likely set of sectors that may be included in the Active Set. The serving RNC then corresponds directly with the RN's where these sectors reside, to request traffic channel and backhaul resources. The RN's either decline, or accept and allocate the needed radio resources. If resources are available from a sufficient set of RN's, the serving RNC accepts the connection request, and sends a TrafficChannel assignment message over the Control Channel to the AT. The AT then starts transmitting on the Reverse Traffic Channel (RTC). Once it acquires the RTC, the RN sends an RTCAck message to the AT to indicate the acquisition of the RTC signal. The AT then responds with a TrafficChannelComplete message to indicate the completion of the Connection set-up.

In this procedure each RN controls its own radio resources, both with respect to hardware resources available on the RN, as well as the management of interference across its sectors. As a result, the admission control function is split between the RN and the serving RNC. RN's provide local admission control for the sectors they control while the serving RNC provides a global admission control. Similarly, when a sector in a given connection is inactive for some period of time, it can initiate the procedure for closing the connection by sending a request to the serving RNC to close the connection. The serving RNC then makes a global decision on whether to remove that sector from the connection, close the entire connection or do nothing.

Packet Routing Between RN and RNC - In More Detail

When a sector in the RN receives a MAC Layer packet on a reverse traffic channel, it forwards the packet to an I/O card after adding a Stream Identifier that includes the UATI of the sending AT along with its own SectorID. The I/O card uses the UATI
5 value to look up the IP address of the serving RNC. It then encapsulates the MAC Layer packet together with its Stream Identifier in an IP packet whose destination address is set to the IP Address of the serving RNC. The I/O module in the serving RNC, upon receiving the packet, reads the UATI value to determine the
10 server module that handles this session. It then passes the packet along with the Stream Identifier to that server module for further processing.

When a sector in the RN receives a MAC Layer packet on the access channel, it first reads the UATI in the ATI field of the MAC
15 Layer Header and then forwards the packet to an I/O card after adding a Stream Identifier that includes the UATI of the sending AT along with the serving sector's SectorID. The I/O card in the RN again uses the UATI value to look up the IP address of the serving RNC. It encapsulates the MAC Layer packet together with
20 its Stream Identifier in an IP packet whose destination address is set to the IP Address of the serving RNC. The I/O module in the serving RNC, upon receiving the packet, reads the UATI value to determine the server module that serves this session. It then passes the MAC Layer packet along with the Stream Identifier to that
25 server module for further processing.

When a server module in the serving RNC has a MAC Layer packet ready for transmission on a forward traffic channel, it first

Attorney Docket 12144-007001

sends it to the I/O card in the serving RNC along with a Stream Identifier that includes the transmitting sector's SectorID (or a representation of it), the UATI of the receiving AT and a MAC Index identifying the connection. The I/O card in the serving RNC

5 then uses the UATI value to look up the IP address of the RN to which to send the packet. It encapsulates the MAC Layer packet together with a Stream Identifier in an IP packet whose destination address is set to the IP Address of the RN. The RN, upon receiving the packet, reads the SectorID value in the Stream

10 Identifier to determine the sector that will transmit the packet. It then passes the MAC Layer packet along with the Stream Identifier to the appropriate modem card, which schedules the MAC Layer packet for transmission on the Forward Link using the MAC Index as the preamble.

15 Similarly, on the forward link, when a server module in the serving RNC has a MAC Layer packet ready for transmission on the Control Channel of a particular sector, it first sends the packet to an I/O card in the serving RNC along with a Stream Identifier that includes the UATI of the receiving AT, the transmitting sector's

20 SectorID (or a representation of it) and a MAC Index identifying the packet as a control channel packet. The I/O card in the serving RNC then uses the UATI value to determine the IP address of the RN to which to send the packet. It then encapsulates the MAC Layer packet together with its Stream Identifier in an IP packet

25 whose destination address is set to the IP Address of the RN. The RN, upon receiving the packet, reads the SectorID value in the Stream Identifier to determine the sector that will transmit the

Attorney Docket 12144-007001

packet. It then passes the MAC Layer packet along with the SectorID and MAC Index to the appropriate modem card. The modem card schedules the packet for transmission on the control channel.

5 **FAILURE RECOVERY & LOAD BALANCING**

The client/server architecture described earlier can be further extended to increase the overall reliability of the wireless network. (Note, the RNC may be a carrier-class equipment with internal redundancy to handle failure of its various cards/modules. The
10 situation we consider here is one where the equipment either does not have redundancy for every card/module or where the redundant component also fails.)

Failure Recovery Without Session Preservation

First, consider an approach where each RN, upon power-up, first
15 communicates with a primary RNC Resource Control Agent who may reside in one or more of the RNC's. The primary Resource Control Agent assigns each RN to a Default RNC. The RN then routes all new session requests to that Default RNC.

When an RNC becomes completely unreachable due to some
20 failure, all AT's that are being served by that RNC will ultimately recognize that their IS-856 sessions have been lost. Each of these AT's will initiate a new session by sending a UATI_Request over the Access Channel. Every RN who receives one of these requests will route them to its default RNC. If at any time, the RN cannot
25 reach its default RNC, it will immediately request a new default

Attorney Docket 12144-007001

RNC from the primary RNC Resource Control Agent. If the primary RNC Resource Control Agent is also not reachable, it will send a similar request to a secondary RNC Resource Control Agent. Once the UATI_Request is received by the Default RNC, it
5 will immediately establish a new IS-856 session with the AT and will further initiate the procedure to set up a new A10 connection with a PDSN.

Assignment of a new Default RNC may also be initiated by the RNC Resource Control Agent. This can be accomplished by
10 having the RNC Resource Control Agent continuously monitor the health of all the RNC's in the subnetwork. Upon detecting the failure of an RNC, the RNC Resource Control Agent immediately communicates with all affected RN's and assigns them to new Default RNC's. In assigning RN's to Default RNC's, the RNC
15 Resource Control Agent may perform load balancing to ensure that user sessions are evenly distributed across all available RNC's.

Load Balancing Session Assignment

The above method can be further enhanced by making the RNC Resource Control Agent ultimately responsible for assigning user
20 sessions to RNC's. In this case, when a Default RNC or possibly the RN itself receives a new UATI_Request, it asks the RNC Resource Control Agent to assign the session to an RNC. The RNC Resource Control Agent assigns the session to an RNC based on resource availability, loading and the distance between the RNC
25 and the RN presently serving the AT. This approach provides better load balancing among RNC's, allowing user sessions to be

distributed across RNC's more dynamically, while also taking into account the current position of the AT. In case of an RNC failure, all new session requests will arrive at the RNC Resource Control Agent who will then assign these sessions to new RNC's, again
5 based on loading and other considerations.

The RNC Resource Control Agent may also be used to trigger dormant handoffs for load balancing or other purposes. In Phase 1 IS-856 networks, a dormant inter-RNC handoff is always triggered by the AT upon detection of a subnet change. As we discussed
10 earlier, lack of an immediate dormant handoff may result in lost paging data.

In the improved IS-856 networks shown in Figures 3 and 4, a dormant handoff can be initiated by the network based on the location of the AT. Upon receipt of a RouteUpdate, when a
15 serving RNC determines that a transfer of a user session to another RNC is desired (for load balancing or other reasons), it sends a Dormant Handoff request to the RNC Resource Control Agent who assigns the session to a new RNC. The new serving RNC then assigns a new UATI and performs a session transfer from the
20 previous serving RNC.

In a more distributed implementation of the RNC Resource Control Agent concept, RNC's can constantly communicate with the RN's and other RNC's to provide routing information (including their loading) to all the RN's, thereby allowing the RN's
25 to route incoming session requests to the correct RNC without going through a RNC Resource Control Agent. A drawback of this

approach is that significant backhaul signaling traffic would be created as a result of exchanging such dynamic loading information.

Failure Recovery with Session Preservation

- 5 In some networks, it may be necessary to recover user session information in case of an RNC failure. This would eliminate the air link congestion that hundreds of new session requests could create shortly after an RNC failure. In order to preserve sessions in case of failure of an RNC, a copy of such information (for all
10 sessions in the subnetwork) can be stored in the RNC Resource Control Agent.

- When an RNC fails and the AT initiates a new session, its new session request will reach the RNC Resource Control Agent. The RNC Resource Control Agent then not only assigns a new serving
15 RNC to each session, but also provides the session information thereby avoiding lengthy session establishment procedures. Once a new UATI is successfully assigned to the AT, communication with the network may resume. The RNC Resource Control Agent further provides information related to the A10 interface, in order
20 to allow the RNC establish an A10 session with the same PDSN, thereby avoiding the setting up of new PPP and Mobile/Simple IP sessions.

- A similar recovery procedure can be applied within the RNC, by setting up an RNC Resource Control Agent inside the RNC. The
25 RNC Resource Control Agent may then run on a specific redundant card, with a hot standby. The RNC Resource Control

Attorney Docket 12144-007001

Agent is then responsible for allocating sessions to server modules. In case a server module fails, the session is internally reallocated to another server module. In principle, the operation of this system is the same as the one operating across the network. Moreover, in
5 this case, it is not necessary to reestablish the A10 session to the PDSN, since the external IP address of the PCF seen by the PDSN can be maintained.

INTEGRATED RNC & PDSN

Another benefit of the client-server architecture described above is
10 the ability to combine the RNC and PDSN functions in a single network element. In hierarchical 3G packet data networks, a PDSN represents the highest point in the hierarchy, and therefore can support multiple RNC's. A new generation of PDSN's are expected to supports hundreds of thousands of users, and several
15 RNC's.

In existing radio access networks with dedicated point-to-point links between RN's and RNC's, migrating the PDSN function to the RNC would be undesirable, because this would reduce the number of sessions that could be supported, resulting in frequent
20 costly handoffs between PDSN's that involve new PPP and Simple/Mobile IP registrations.

In the client/server architecture described here handoffs between RNC's occur much less frequently therefore allowing the integration of the PDSN function into the RNC. Such an approach
25 also simplifies the networking between the RNC and the PDSN, and further increases scalability and reliability.

Attorney Docket 12144-007001

In an RNC with an integrated PDSN, so-called PDSN server modules are added to handle the PDSN function. This includes PPP termination, Simple IP and/or Mobile IP foreign agent and the AAA client functions. As long as the AT remains within a subnet
5 (say an RNC cluster), no inter-PDSN handoffs would be required.

If an integrated RNC/PDSN fails, all sessions supporting an AT (including the air interface, PPP and Simple/MobileIP sessions) are transferred to another RNC/PDSN thereby avoiding any new session establishment between the AT and the wireless network.

10 Other embodiments are within the scope of the following claims.